

Supplementary Materials

Appendix A. Task-specific prediction performance on HCP dataset

In Table 2, we only reported the averaged performance over tasks. We stated that BrainGNN significantly outperformed the alternative methods. To have more insight into how BrainGNN performs on individual tasks, we report the mean and standard deviation of the metrics on each task over the testing sets.

Table A.3: Task-wise classification performance using BrainGNN classifier.

	Gambling	Language	Motor	Relational	Social	WM	Emotion
Accuracy(%)	87.86(3.42)	96.28(2.24)	96.08(1.98)	91.18(3.46)	96.88(1.45)	93.52(1.11)	93.54(1.76)
F1(%)	90.20(2.28)	97.20(1.10)	96.20(1.70)	91.80(2.17)	98.80(0.84)	93.60(1.67)	93.80(1.30)
Recall(%)	88.40(3.36)	97.00(1.87)	96.80(1.79)	91.80(3.63)	97.60(1.14)	94.20(0.84)	94.20(1.79)
Precision(%)	91.40(2.41)	97.20(1.30)	95.20(2.39)	92.00(4.18)	99.20(0.84)	93.00(3.08)	92.80(2.28)

Appendix B. Justification on the number of communities $K = 8$

Intuitively, a small number of communities K will reduce the learnability of the model, while a large K increases the learnable parameters. We are motivated by Finn et al. (2015) and set the number of communities $K = 8$. To show $K = 8$ is a reasonable selection, we report the dice score (Zou et al., 2004) of the top salient ROIs selected by the 1st R-pool layer in each community in Fig. B.11, which measures the overlap of the saliency areas with each community. The results were from the best fold of each dataset. The R-pool layer tends to select a few representative ROIs in each community to summarize the group-level representation. We also noticed that a node within multiple communities tends to be selected as well.




Fig. B.11: The dice scores of the overlap between the selected salient areas and each community. Each community has a small portion of ROIs being saliency to a specific prediction.

Appendix C. Visualize the important interactions between the important ROIs

Functional connectivity has been used as important brain biomarkers in many studies (Venkataraman et al., 2016; Ozdemir et al., 2020; Gao et al., 2020). In this paper, we focus on extracting patterns from nodes' representations. The edge connections \tilde{e}_{ij} served as message passing filters (see Eq. 1). Finding the best measurement of the importance of functional connectivity is often ambiguous due to the intrinsic complex of fMRI. Without the help of additional post-hoc interpretation methods, we can infer the connections between the important nodes as the important functional connections, as the results show in Fig. C.12 and C.13. If we want to filter out the important functional connections through edge embedding, this future work can be generalized from our model, such as by applying edge-convolution (Wang et al., 2019b) and edge-pooling (Diehl et al., 2019; Diehl, 2019).




Fig. C.12: Interpreting important connectivity of Biopoint task. The connections among the top salient ROIs shown in Fig 6.

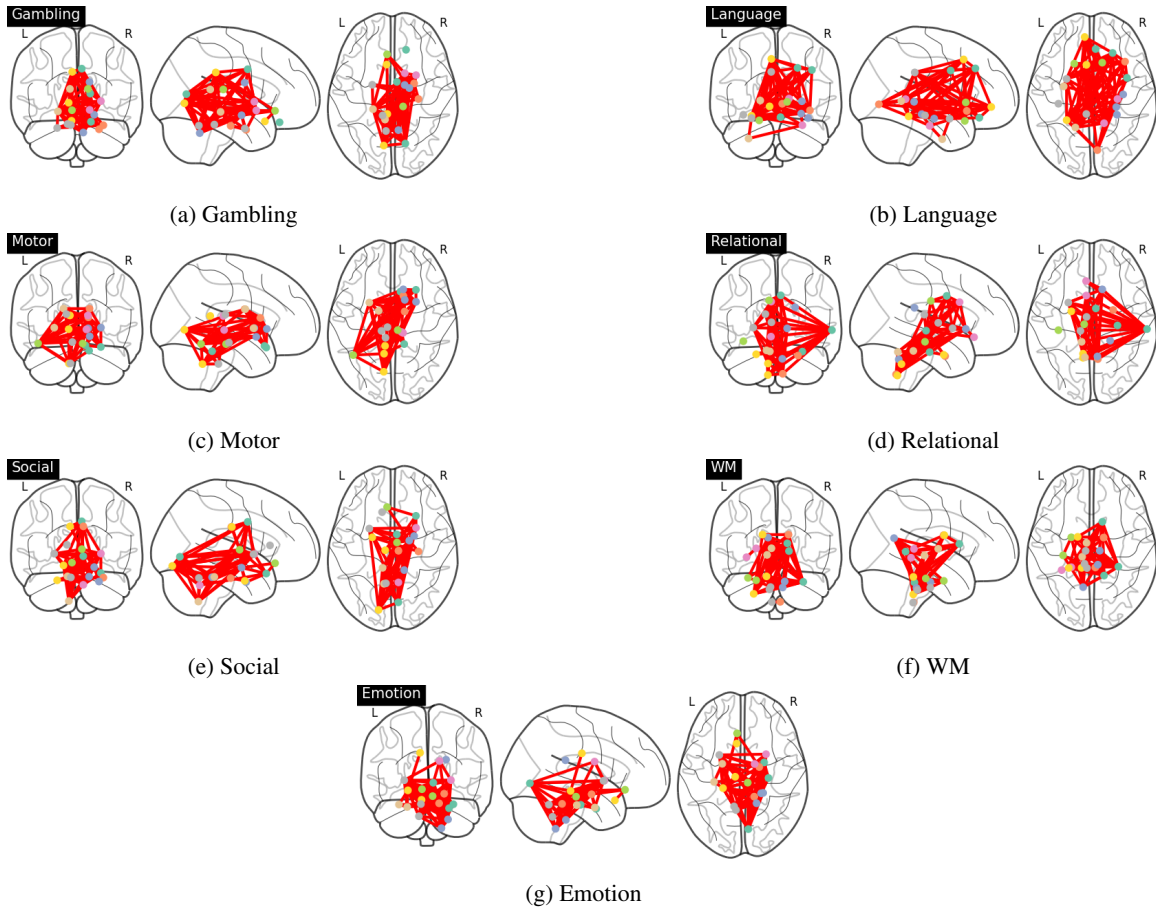


Fig. C.13: Interpreting important connectivity of HCP task. The connections among the top salient ROIs shown in Fig 7.